AMENDMENTS TO THE SPECIFICATION

Please amend the following paragraphs beginning at page 8, line 23:

Referring to FIG. 1, it ~~It~~ is to be understood that the present invention may be implemented in various forms of hardware, software, firmware, special purpose processors, or a combination thereof. In one embodiment, the present invention may be implemented in software as an application program tangibly embodied on a program storage device. The application program may be uploaded to, and executed by, a machine 101 comprising any suitable architecture. Preferably, the machine 101 is implemented on a computer platform having hardware such as one or more central processing units (CPU) 102, a random access memory (RAM) 103 and input/output (I/O) interface(s) 104. The computer platform also includes an operating system and microinstruction code. The various processes and functions described herein may either be part of the microinstruction code or part of the application program (or a combination thereof) 107, which is executed via the operating system. In addition, various other peripheral devices 105/106/108 may be connected to the computer platform 101 such as an additional data storage device and a printing device.

It is to be further understood that, because some of the constituent system components and method steps depicted in the accompanying figures may be implemented in software, the actual connections between the system components (or the process steps) may differ depending upon the manner in which the present invention is programmed. Given the teachings of the present invention provided herein, one of ordinary skill in the related art will be able to contemplate these and similar implementations or configurations of the present invention.

Please amend the paragraph beginning at page 14, line 16:

The residual image motion after plane alignment obeys certain properties. Referring to FIG. 2, by ~~By~~ way of example, let P be a point not on the plane that is registered, and let p be its image in a reference view. Let T_1 denote the baseline vector between the cameras and Q be the point where the ray connecting P to the second camera center intersects the surface. Then the residual parallax displacement ~~cau~~ u at image location p can be shown to be

$$\delta\mathbf{u} = \mathbf{q} - \mathbf{p} = \frac{T_z(Q_z - P_z)}{Q_z(P_z - Tz)}(\mathbf{p} - \mathbf{t}_1)$$

where P_z and Q_z denote the depths of points P and Q, T_z is the z component of translation vector T_1, and t_1 denotes the epipole corresponding to T_1. If the surface that is aligned is a plane, then the residual parallax displacement simplifies in the case of $T_z \gamma 0$ to:

$$\delta u = \frac{HT_z(\mathbf{p} - \mathbf{t}_1)}{T_\perp P_z - HT_z} = \frac{HT_z}{T_\perp P_z}(\mathbf{q} - \mathbf{t}_1) \qquad (2)$$

and in the case of $T_z = 0$ to:

$$\delta u = -\frac{fH}{T_\perp P_z}\vec{T}_1 \qquad (3)$$

where H is the perpendicular distance from the point P to the reference plane, T_ζ is the perpendicular distance between the second camera center M and the reference plane.

Please amend the paragraph beginning at page 16, line 11:

After determining the parallax between the two views (for example, see FIGs. 3a and

3b), changes in the parallax space are identified. Parallax at a particular point is dependent only on the structure of the scene and is insensitive to changes in illumination, shadows, specularities etc. Therefore, it is more robust than intensity based methods. Objects can be identified as regions with a different parallax than the background. This can be further combined with intensity-based methods to obtain detection. According to an embodiment of the present invention, parallax is used with uncalibrated PTZ cameras. A probability density function (pdf) is determined for the observed parallax (e.g., along with intensity) at a pixel. Such density may be estimated by several techniques such as the mixture-of-Gaussians or Non-parametric kernels.

Please amend the following paragraphs beginning at page 18, line 15:

In these equations, T_z and T_ζ are the same for all points in the scene, while H and P_z vary. Moreover, given the point match, the point q and disparity δu are available. One can perform a simple calibration of the scene by providing the real locations of four ground plane points. This information need not be very accurate and an approximate estimate is sufficient. This will give us a homography between the actual ground plane and its image in the camera (see for example, FIG 4). Using this information, the distance P_z of an observed point can be approximated if the distance of the object from the ground plane is assumed to be small. P_z is determined up to a global scale factor. It is possible to determine such scale factor automatically from the scene by tracking objects and observing their changes in the disparity as they move about. (see FIG. 5 showing a disparity map between FIGs. 3a and 4). This can be performed if it is assumed that the object does not change its shape drastically during the motion. This is a reasonable assumption for moving cars and people.

Thus, apart from an unknown global scale factor T_z, which can be estimated if the exact location of the second camera is available, one can correct for other alterations in the parallax as an object moves in the scene (see for example, FIG. 6). Hence, for a particular object having a non-altering height (for example, cars or walking people), one can obtain a quantity that remains unaltered as the object moves around the scene. This information can be used to track an object as it moves in the scene and to recover from occlusions. Furthermore, such representation allows us to recognize/classify an object since the structure of the object is captured in this invariant space.

Please amend the following paragraphs beginning at page 25, line 2:

According to an embodiment of the present invention and referring to FIG. 7, a method of tracking an object comprises providing at least two cameras 701. The cameras can be uncalibrated. The method further comprises determining an image from each camera 702, and determining a common plane in the acquired images 703. The method further comprises determining parallax for all points in the images 704, and incorporating parallax in a background model 705. The background model can comprise other features such as intensity and edges. The method comprises estimating a change in the scene in the parallax+background joint space 706.